

# Automatic TM Cleaning through MT and POS Tagging: Autodesk's Submission to the NLP4TM 2016 Shared Task

Alena Zwahlen, Olivier Carnal, Samuel Lübli

Autodesk Development S.à.r.l.

Rue du Puits-Godet 6, 2000 Neuchâtel, Switzerland

{alena.zwahlen, olivier.carnal, samuel.laubli}@autodesk.com

## Abstract

We describe a machine learning based method to identify incorrect entries in translation memories. It extends previous work by Barbu (2015) through incorporating recall-based machine translation and part-of-speech-tagging features. Our system ranked first in the Binary Classification (II) task for two out of three language pairs: English–Italian and English–Spanish.

**Keywords:** Translation Memory, Machine Learning, Machine Translation, Part-of-Speech Tagging

## 1. Introduction

Autodesk has accumulated more than 40 million professionally translated segments over the past 17 years. These translation units (TUs) mainly stem from user interfaces and documentation of software products localized into 32 languages. As we are now unifying and centralizing all translations in a single repository, it is high time to sort out duplicate, outdated, and erroneous TUs. Exploring methods to handle the latter – clearly more challenging than removing duplicate and outdated material – motivated us to participate in the First Shared Task on Translation Memory Cleaning (Barbu et al., 2016). Going forward, we strive to make human translation more efficient (by showing translators less erroneous fuzzy matches) and machine translation more accurate (by reducing noise in training data).

In this paper, we describe our submitted system for distinguishing correct from incorrect TUs. Rather than tailoring it to individual languages, we aimed at a language-independent solution to cover all of the language pairs in this shared task or, looking to the future, Autodesk's production environments. The system is based on previous work by Barbu (2015) and uses language-independent features with language-specific plug-ins, such as machine translation, part-of-speech tagging, and language classification.

Specifics about previous work are given in the next section. In Section 3, we describe our method and, in Section 4, show how it compares to Barbu's (2015) approach as well as other submissions to this shared task. Lastly, we offer preliminary conclusions and outline future work in Section 5.

## 2. Background

TM cleaning functionality in commercial tools is mostly rule-based, centering around the removal of duplicate entries, ensuring markup validity (e.g., no unclosed tags), or controlling for client or project specific terminology. Although helpful, these methods fall short of identifying spurious entries that contain language errors or partial translations. With crowd-sourced and automatically constructed TMs in particular, it is also necessary to identify transla-

tion units with source and target segments that do not correspond at all (e.g., Trombetti, 2009; Tiedemann, 2012).

Barbu (2015) has proposed to cast the identification of such incorrect translations as a supervised classification problem. In his work, 1,243 labelled TUs were used to train binary classifiers based on 17 features. The “most important” of them, according to the author, were *bisegment\_similarity* and *lang\_diff*: the former is defined as the cosine similarity between a target segment and its machine translated source segment, while the latter denotes whether the language codes declared in a translation unit correspond with the codes detected by a language detector. The best classifier, a support vector machine with linear kernel, achieved 82% precision and 81% recall on a held-out test set of 309 TUs.

To the best of our knowledge, Barbu provided the first and so far only research contribution on automatic TM cleaning, which the author himself described as “a neglected research area” (Barbu, 2015). With our participation to this shared task, we seek to extend his work by examining new features based on statistical MT and POS tagging.

As outlined above, comparing machine translated source segments to their actual target segments has proven effective in Barbu's (2015) experiments. We propose to complement or replace the similarity function used for this comparison (cosine similarity) by two automatic MT evaluation metrics, Bleu (Papineni et al., 2002) and character-based Levenshtein distance, in order to reward higher-order  $n$ -gram ( $n > 1$ ) and partial word overlaps, respectively. Furthermore, we introduce a recall-based MT feature that takes multiple MT hypotheses ( $n$ -best translations) of a given source segment into account, based on the assumption that alternative translations of words (such as “buy” and “purchase”) or phrases (such as “despite” and “in spite of”) should not be punished.

We also experiment with part-of-speech information to identify spurious translation units. With closely related languages in particular, the rationale would be that adjectives (to name an example) in a source segment are likely to be reflected in the corresponding target segment in case of a valid translation. The comparison of POS tags from language-specific tagsets will be based on a mapping to

Language	Domain	Set	Translation Units			
			1	2	3	Total
en-de	News	Training	1,086	100	210	1,396
		Evaluation	544	51	105	700
en-es	Medical	Training	942	128	313	1,383
		Evaluation	471	65	157	693
en-it	Medical	Training	872	254	284	1,410
		Evaluation	437	128	143	708

Table 1: Trainig and evaluation data. Classes 1, 2, and 3 denote correct, almost correct, and incorrect translation units, respectively.

eleven coarse-grained, language-independent grammatical groups (Petrov et al., 2011).

We acknowledge that the use of MT is discouraged by the organizers of this shared task to foster contributions that require less compute power. However, as MT was found to be valuable in previous work (see above) and computational resources are hardly a limiting factor in corporate environments (see Section 3.2), we decided not to refrain from including MT-based features in our submissions.

### 3. Method

Our system uses labelled TUs to train classifiers based on language-independent features (see Section 3.1) with language-specific plug-ins (see Section 3.2). The feature extraction pipeline is implemented in Scala (see Section 3.3), and our final submission – geared to distinguish correct or almost correct from incorrect TUs – is based on a selection of nine features (see Section 3.4).

#### 3.1. Features

We re-implemented the 17 features proposed by Barbu (2015, see also Section 2). In addition, we explore

- *mt\_coverage* the percentage of target words contained in the  $n$ -best machine translations of the source segment. We use  $n = 20$  in our experiments.
- *mt\_cfs* the character-based Levenshtein distance between target segment and machine translated source segment. We normalise this score such that identical and completely dissimilar segments result in scores of 1.0 and 0.0 respectively, i.e.,

$$cfs = 1 - \frac{\text{Levenshtein distance in characters}}{\text{number of characters in longer segment}}.$$

This score is computed individually for each of the 20-best translation options; the best of these scores instantiates the feature value.

- *mt\_bleu* the BLEU score (Papineni et al., 2002) between target segment and machine translated source segment. We employ the sentence-level version of the metric as implemented in Phrasal (Green et al., 2014). As with *mt\_cfs*, individual scores are computed for each of the 20-best translation options; the best score instantiates the feature value.

- *pos\_sim\_all* the cosine similarity between the part-of-speech (POS) tags found in the source and target segment.
- *pos\_sim\_some* the cosine similarity between source and target segment in terms of nouns (NOUN), verbs (VERB), adjectives (ADJ), and pronouns (PRON).
- *pos\_exact* whether or not the POS tag sequence in source and target segment is identical.
- *language\_detection* whether or not a state-of-the-art language classifier confirms the target segment’s language declared in the translation unit.
- *ratio\_words* the ratio between number of words in source and target segment.
- *ratio\_chars* the ratio between number of characters in source and target segment.

#### 3.2. Resources

Some of the features described in the previous section require natural language processing (NLP) facilities. For machine translation, we use our in-house systems (Plitt and Masselot, 2010; Zhechev, 2014) based on the Moses SMT framework (Koehn et al., 2007). They are trained on translated software and user manuals from Autodesk products only and chosen for the sake of convenience; we would expect better performance of our MT-based features in conjunction with MT engines geared to the text domains used in this shared task (listed in Table 1). Our engines are integrated into a scalable infrastructure deployed on an elastic compute cloud, allowing high throughput even with large translation memories to be cleaned.

For POS tagging, we rely on Schmid’s (1995) TreeTagger and its readily available models<sup>1</sup> for English, German, Italian, and Spanish. To make POS tags comparable across these languages, they are mapped<sup>2</sup> to the Universal Tagset proposed by Petrov et al. (2011).

Lastly, we use the publicly available Xerox Language Identifier API<sup>3</sup> for language detection.

#### 3.3. Classification

Our feature extraction pipeline, including Barbu’s (2015) as well as our own features (see Section 3.1), is implemented in Scala. This pipeline is used to transform translation units into feature vectors and train classifiers using the scikit-learn framework (Pedregosa et al., 2011). From the various classification algorithms we tested, Random Forests performed best with our selection of features (see below).

#### 3.4. Feature Selection

For the reasons mentioned in Section 1, we aimed at finding a combination of features that would perform well with all language pairs rather than tailoring solutions to individual

<sup>1</sup><http://www.cis.uni-muenchen.de/~schmid/tools/TreeTagger/>

<sup>2</sup><https://github.com/slavpetrov/universal-pos-tags>

<sup>3</sup><https://open.xerox.com/Services/LanguageIdentifier>

Features	Algorithm	en-de			en-es			en-it		
		P	R	F <sub>1</sub>	P	R	F <sub>1</sub>	P	R	F <sub>1</sub>
Baseline 1	Random + Class Distribution	.77	.77	.77	.67	.68	.67	.70	.71	.70
Baseline 2	Adapted Church-Gale	.78	.77	.77	.72	.71	.71	.76	.74	.74
Barbu (2015)	SVM (linear kernel)	.74	.85	.78	.84	.85	.83	.83	.83	.79
Barbu (2015)	Random Forests	<b>.80</b>	.85	<b>.79</b>	.88	.88	.87	.87	.88	.87
This work	Random Forests	<b>.80</b>	<b>.86</b>	<b>.79</b>	<b>.89</b>	<b>.89</b>	<b>.89</b>	<b>.90</b>	<b>.90</b>	<b>.90</b>

(a) Binary Classification (II) task.

Features	Algorithm	en-de			en-es			en-it		
		P	R	F <sub>1</sub>	P	R	F <sub>1</sub>	P	R	F <sub>1</sub>
Baseline 1	Random + Class Distribution	.63	.63	.63	.53	.54	.53	.45	.45	.45
Baseline 2	Adapted Church-Gale	.64	.63	.63	.57	.57	.57	.49	.48	.48
Barbu (2015)	SVM (linear kernel)	.63	.78	.69	.70	.77	.73	.68	.70	.61
Barbu (2015)	Random Forests	.72	.78	.70	<b>.80</b>	<b>.81</b>	<b>.79</b>	<b>.75</b>	<b>.75</b>	<b>.72</b>
This work	Random Forests	<b>.77</b>	<b>.78</b>	<b>.70</b>	.74	<b>.81</b>	.76	<b>.75</b>	<b>.75</b>	.70

(b) Fine-Grained Classification task.

Table 2: Classification results on training data.

Team	en-de		en-es		en-it	
	F <sub>1</sub>	c	F <sub>1</sub>	c	F <sub>1</sub>	c
Baseline 1	.53	534	.43	418	.50	473
Baseline 2	.52	532	.51	442	.56	492
Autodesk	.47	593	<b>.81</b>	<b>611</b>	<b>.85</b>	<b>644</b>
Christian Buck-MT	.66	597	–	–	–	–
Christian Buck-NMT	.65	594	–	–	–	–
FBK-HLTMT	.49	594	.77	596	.80	631
JUMTTeam	.58	482	.66	493	.70	530
Lingua Custodia	.64	609	.78	592	.83	634
Unisa	<b>.68</b>	<b>618</b>	.76	596	.77	623

(a) Binary Classification (II) task: averaged F<sub>1</sub>-scores and number of correctly classified TUs (c).

Team	en-de	en-es	en-it
	F <sub>1</sub>	F <sub>1</sub>	F <sub>1</sub>
Baseline 1	.64	.48	.47
Baseline 2	.63	.52	.50
Autodesk	.69	.74	.68
Christian Buck-MT	.77	–	–
Christian Buck-NMT	.77	–	–
FBK-HLTMT	.70	.72	.66
Lingua Custodia	.75	<b>.79</b>	.69
Unisa	<b>.80</b>	.77	<b>.73</b>

(b) Fine-Grained Classification task: weighted F<sub>1</sub>-scores.

Table 3: Official evaluation and ranking results. Winning systems are highlighted in bold type.

languages. We focused on gearing our classifiers to distinguish correct or almost correct (classes 1, 2) from incorrect TUs (class 3) – i.e., the Binary Classification (II) task – by optimising the weighted F<sub>1</sub>-score (F<sub>1</sub>) on training data (see Tables 2a and 2b). From the various feature combinations we tested, we found the following to be most successful: *ratio\_words*, *pos\_sim\_all*, *language\_detection*, *mt\_cfs*, *mt\_bleu*, *ratio\_chars* (as described in Section 3.1), alongside *cg\_score*, *only\_capletters\_dif*, and *punctuation\_similarity* (from Barbu, 2015). Evaluation results are given in the next section.

## 4. Results

We tested our final submission – a Random Forests classifier based on the nine features described in Section 3.4 – on three language pairs (en-de, en-es, en-it) and two tasks: Binary II and Fine-Grained Classification (see Sections 4.1 and 4.2, respectively). The classifier was trained solely on data provided by the organizers of this shared task for each of the language × task conditions. Each TU in this data was annotated with one of three labels: correct, almost correct, and incorrect (see Table 1).

### 4.1. Binary Classification (II)

Our rationale for focusing on telling apart correct or almost correct from incorrect TUs was that a first application of our method, if successful, would most likely be the filtering of TM data for MT training. While eliminating almost correct TUs might decrease rather than increase MT quality, filtering out incorrect segments can have a positive impact (Vogel, 2003).

Prior to submission, we benchmarked our system against the two baselines provided by the organizers: a dummy classifier assigning random classes according to the overall class distribution in the training data (Baseline 1), and a classifier based on the Church-Gale algorithm as adapted by Barbu (2015) (Baseline 2). More importantly, however, we compared our system to Barbu’s (2015) approach, using the classification algorithms which reportedly worked best with the 17 features in his work. Our system performed well in this comparison, surpassing Barbu’s approach in all language pairs except en-de, where both systems were *en par*. Details are shown in Table 2a, where we report weighted precision (P), recall (R), and F<sub>1</sub>-scores averaged over 5-fold cross-validation with  $\frac{2}{3}$ – $\frac{1}{3}$  splits of the training data.

The final evaluation and ranking produced by the organizers, shown in Table 3a, confirms our findings from experimenting with training data: our system performs well on the en-es and en-it test sets (best in class), while performance is substantially lower on the en-de test set. The reasons for this are yet to be ascertained (see also Section 5).

## 4.2. Fine-Grained Classification

Although geared to the Binary Classification (II) task (see above), we also assessed our system on the Fine-Grained Classification task. Here, the goal was to distinguish between all of the three classes, i.e., determine whether a TU is correct, almost correct, or incorrect.

Again, we compared our system’s performance to Barbu’s (2015) method, using  $\frac{2}{3}$ - $\frac{1}{3}$  splits of the training data (5-fold cross-validation). The results, shown in Table 2b, implied that the nine features we selected would not suffice for a more fine-grained classification of TUs. This was confirmed in the official evaluation and ranking: our system scored low on en-de and mediocre on en-es and en-it. Further work will be needed to analyse these results in more detail.

## 5. Conclusions

We have proposed a machine learning based method to identify incorrect entries in translation memories. It is applicable to any language pair for which an MT system, a POS tagger, and a language identifier are available. Implemented using off-the-shelf tools, our system achieved the best classification results for two out of three language pairs (English-Italian and English-Spanish) in the Binary Classification (II) task.

In future work, we would like to assess the impact of gearing NLP components to target domains on classification accuracy. The training data in this shared task stems from news (German) and medical texts (Italian, Spanish) which our MT systems, for example, were not optimized for. This domain mismatch might partially explain why our system did not perform well on the English-German test set.

More importantly, however, we would like to test our implementation as-is in Autodesk’s production environments for software localization. Removing incorrect segments from TMs could ultimately help make professional translation more efficient by providing better MT (through filtered training data) and more accurate fuzzy matches.

## 6. Acknowledgements

We would like to thank Valéry Jacot for his vital support and guidance.

## 7. Bibliographical References

Barbu, E. (2015). Spotting false translation segments in translation memories. In *Proceedings of the Workshop on Natural Language Processing for Translation Memories (NLP4TM)*, pages 9–16, Hissar, Bulgaria.

Barbu, E., Escartín, C. P., Bentivogli, L., Negri, M., Turchi, M., Federico, M., Mastrostefano, L., and Orasan, C. (2016). 1st shared task on automatic translation memory cleaning preparation and lessons learned. In *Proceedings*

*of the 2nd Workshop on Natural Language Processing for Translation Memories (NLP4TM)*, Portorož, Slovenia.

Green, S., Cer, D., and Manning, C. D. (2014). Phrasal: A toolkit for new directions in statistical machine translation. In *Proceedings of the 9th Workshop on Statistical Machine Translation (WMT)*, pages 114–121, Baltimore, USA.

Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R., Dyer, C., Bojar, O., Constantin, A., and Herbst, E. (2007). Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 177–180, Prague, Czech Republic.

Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (2002). Bleu: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 311–318, Philadelphia, Pennsylvania.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Petrov, S., Das, D., and McDonald, R. (2011). A universal part-of-speech tagset. *arXiv preprint*, arXiv:1104.2086.

Plitt, M. and Masselot, F. (2010). A productivity test of statistical machine translation post-editing in a typical localisation context. *Prague Bulletin of Mathematical Linguistics*, 93:7–16.

Schmid, H. (1995). Improvements in part-of-speech tagging with an application to German. In *Proceedings of the ACL SIGDAT Workshop*, pages 47–50, Dublin, Ireland.

Tiedemann, J. (2012). Parallel data, tools and interfaces in OPUS. In *Proceedings of the 8th International Conference on Language Resources and Evaluation (LREC)*, pages 2214–2218, Istanbul, Turkey.

Trombetti, M. (2009). Creating the world’s largest translation memory. In *Proceedings of the 12th Machine Translation Summit (MT Summit)*, Ottawa, Canada.

Vogel, S. (2003). Using noisy biligual data for statistical machine translation. In *Proceedings of Meeting of the 10th Conference of the European Chapter of the Association of Computational Linguistics (EACL)*, pages 175–178, Budapest, Hungary.

Zhechev, V. (2014). Analysing the post-editing of machine translation at Autodesk. In O’Brian, S., Balling, L. W., Carl, M., Simard, M., and Specia, L., editors, *Post-editing of Machine Translation: Processes and Applications*, pages 2–23. Cambridge Scholars Publishing.